REMARKS

Claims 1-26 were pending in the present application. Claims 1-3, 10-11, and 18-26 were amended. Accordingly, claims 1-26 remain pending in the application.

Claims 3, 20, and 18-25 were objected to for informalities. Applicant has amended claims 3 and 20, and renumbered the claims as requested by the Examiner to overcome the objections.

Claims 1, 4-8, 10, 15-18, and 23-26 stand rejected under 35 U.S.C. 102(e) as being anticipated by Mudgett et al. (U.S. Patent No. 6,775,749) (hereinafter "Mudgett"). Although Applicant respectfully traverses at least portions of this the rejection, Applicant has amended some of the claims to clarify the claim language. Accordingly, Applicant respectfully requests reconsideration in light of the foregoing amendments and the following remarks.

Claim 9 stands rejected under 35 U.S.C. 103(a) as being unpatentable over Mudgett in view of Cypher (U.S. Patent No. 6,629,205) (hereinafter "Cypher"). Applicant respectfully traverses the rejection.

Claims 2-3, 11-14, and 19-22 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Mudgett in view of Van Doren et al. (U.S. Patent No. 6,202,126) (hereinafter "Van Doren"). Although Applicant respectfully traverses at least portions of this the rejection, Applicant has amended some of the claims to clarify the claim language. Accordingly, Applicant respectfully requests reconsideration in light of the foregoing amendments and the following remarks.

Applicant's claim 1 recites

"A cache memory subsystem comprising:
a cache storage configured to store a plurality of cache lines of data;
a scheduler configured to schedule reads and writes of information to said cache storage using a fixed latency pipeline;

wherein in response to scheduling a read request to said cache storage, said scheduler is further configured to cause **an associated write** to said cache storage to occur a fixed number of cycles after said scheduling a read request." (Emphasis Added)

The Examiner asserts Mudgett teaches each and every limitation recited in Applicant's claim 1. Applicant respectfully disagrees with the Examiner's assertion and characterization of Mudgett and the application of Mudgett to Applicant's claims. More particularly, Applicant asserts that although Mudgett deals with processors and cache systems, Mudgett is solving a different problem using a different apparatus and mechanism for doing so.

Specifically, Mudgett is directed to performing speculative cache fills using cache probes and speculative responses, especially in systems including multiple processors where cache coherency is a problem. Mudgett teaches at col. 5, line 42 through col. 6, line 4

> "Each processor 102 may be configured to request data from memory 106 whenever that data is not currently present in that processor's cache. Before the requested data is provided from a shared portion of the memory 106, however, a cache coherency mechanism (e.g., in chipset 104) may need to ascertain whether the other processor 102B has a copy of the requested line in its cache 110B and, if so, whether the other processor 102B has modified its copy of the requested line. Accordingly, chipset 104 may be configured to send a probe to processor 102B that causes processor 102B to search cache 110B for the requested line. As used herein, a "probe" may refer to any communication that causes another device to determine whether it has a copy of the identified data. In some embodiments, a probe may be sent as a packet on an independent communication channel, while in other embodiments, a probe may be communicated as signals asserted on a shared bus. For example, in one embodiment, chipset 104 (or processor 102A) may send a probe packet to processor 102B in response to processor 102A requesting a cache fill from memory 106. In response, processor 102B may generate a response packet that indicates whether a copy of the requested data is present in its cache 110B.
>
> In order to speed up performance, the chipset 104 may be configured to provide a speculative response from memory 106 to processor 102A's cache fill request before the probe response from processor 102B has been received. This way, if processor 102B does not have a copy of the

requested data, the requesting processor 102A may more quickly receive its requested data from system memory 106." (Emphasis added)

Mudgett also discloses at col. 8, line 59 through col. 9, line 13

"A read transaction may be initiated by one of the processors 102A sending a read command to the memory 106 via the chipset 104. When the chipset 104 is ready to return data to processor 102A, a response may be sent to the requesting processor 102A over its bus 112A1 to alert the processor 102A that data is coming on bus 112A3 and to identify the data request to which the response corresponds. The data may have been sent from memory 106 to data buffer 228 on interconnection 114 prior to the chipset sending the response to processor 102A (or alternatively, the data may be provided directly from memory 106 to processor 102A a programmable number of clocks after the chipset 104 generates the response). From data buffer 228, the data may be provided to the processor 102A over the bus 112A3. Similarly, a write transaction may be requested over the requesting processor's processor bus 112A2. The chipset 104 may responsively send a request to the requesting processor 102A that the associated write data be transferred over the memory bus 112A3 to data buffer 228 (or directly to memory 106) a programmable number of clocks later. Chipset 104 may then cause the data to be transferred from data buffer 228 to memory 106 on bus 114." (Emphasis added)

From the foregoing, it is apparent that Mudgett is teaching providing to a requestor, speculative responses (i.e., read response data from system memory) in an effort to reduce latencies associated with waiting for cache probes to be returned from another processor.

In addition, Mudgett is also teaching that in response to a read request to system memory 106, a read response (that may notify the requestor that response data is coming and to which request the data belongs) may be sent to the requestor (one of the processors 102) prior to the read data being sent to the requestor. Mudgett is also teaching that in a similar way, **in response to receiving a write transaction** from a requestor (one of processors 102) the chipset may send a write request back to that processor to request the associated data be sent from the processor to a data buffer in the chipset or to system memory some programmable number of clocks later. Applicant submits *Mudgett's teaching is clearly different* than scheduling a read request to the cache storage and

causing write, which is associated with the read, to the cache storage some fixed number of cycles later in response to scheduling the read request. More particularly, Mudgett does not disclose that the write transaction is associated with the read request or any read for that matter. It is merely an arbitrary write transaction request from the processor.

Furthermore, the pipeline referred to in Mudgett, and pointed out by the Examiner is an instruction pipeline used by the execution units of a processor. Mudgett mentions the pipeline in the context of defining a clock cycle. Thus, Applicant submits Mudgett does not teach scheduling reads and writes to a cache storage in a pipelined fashion.

Van Doren discloses in the abstract "A method for preventing inadvertent invalidation of data elements in a system having a separate probe queue and fill queue for each central processing unit, is provided wherein a central processing unit stores a clean data element, that would otherwise have been discarded, in a victim data buffer when it is evicted from cache."

Applicant submits Van Doren teaches the use of a victim buffer to store cache lines of data evicted from the cache. This is mentioned in Applicant's specification (page 12, lines 5-9) as a problem to be overcome. (i.e., to ensure that victim data will be written to the L2 without the use of a victim buffer).

Cypher is directed to cache memory systems including cache tags that are associated with cache lines of different classes. (See abstract)

Accordingly, Applicant submits neither Mudgett, nor any of the cited references, **teaches or discloses** "schedule reads and writes of information to said cache storage," or "wherein in response to scheduling a read request to said cache storage, said scheduler is further configured to cause **an associated write** to said cache storage to occur a fixed number of cycles after said scheduling a read request" as recited in Applicant's claim 1.

Thus for the reasons given above, Applicant submits claim 1, along with its dependent claims patentably distinguishes over Mudgett.

Applicant's claims 10 and 18 recite language that is similar to the language recited in claim 1. Accordingly, for at least the reasons given above, Applicant submits claims 10 and 18, along with their respective dependent claims, patentably distinguish over the cited art.

CONCLUSION

Applicant submits the application is in condition for allowance, and an early notice to that effect is requested.

If any fees are due, the Commissioner is authorized to charge said fees to Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C. Deposit Account No. 501505/5500-91000/SJC.

Respectfully submitted,

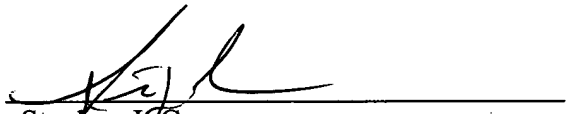

Stephen J. Curran
Reg. No. 50,664
AGENT FOR APPLICANT(S)

Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C.
P.O. Box 398
Austin, TX 78767-0398
Phone: (512) 853-8800

Date: June 12, 2006